Proof Generation from Delta-Decisions

Sicun Gao, Soonho Kong, and Edmund M. Clarke
Carnegie Mellon University, Pittsburgh, PA, USA 15213

arXiv:1409.6414v1 [cs.LO] 23 Sep 2014

***Abstract*—We show how to generate and validate logical proofs of unsatisfiability from delta-complete decision procedures that rely on error-prone numerical algorithms. Solving this problem is important for ensuring correctness of the decision procedures. At the same time, it is a new approach for automated theorem proving over real numbers. We design a first-order calculus, and transform the computational steps of constraint solving into logic proofs, which are then validated using proof-checking algorithms. As an application, we demonstrate how proofs generated from our solver can establish many nonlinear lemmas in the the formal proof of the Kepler Conjecture.**

I. Introduction

Decision solvers for logic formulas over the real numbers play a crucial role in the formal verification of safety-critical embedded systems. For full reliability, decision solvers should provide, besides "sat/unsat" answers, certificates of correctness for such answers. For sat answers, we can certify by just plugging in a solution of the formula (value assignments for all variables). For unsat answers, there is no such witness, and we need mathematical *proofs of unsatisfiability* to guarantee correctness. Such proofs are especially important in the framework of δ-complete decision procedures [4], which rely on numerical procedures that are potentially error-prone. For instance, the following is an actual bug we experienced in building our SMT solver dReal [9]:

> With the standard C library implementation *eglibc-2.15*, included in the latest *Ubuntu 12.10*, the exponential and trigonometric functions contain serious errors. For instance, in upward rounding mode, $\sin(-2.437592) > 10^{53}$. Clearly, this leads to bugs in all constraint solvers using this standard C library.

Note that when we obtain a proof of unsatisfiability, then the correctness of the result becomes independent from the numerical procedures that were used to obtain them.

Besides certifying correctness of solvers, obtaining such proofs is also important from the perspective of automated theorem proving. Decision solvers can establish mathematical theorems by solving satisfiability of the negation of a theorem, and establish correctness through the absence of counterexamples. Valid proofs of unsatisfiability can be directly used as formal proofs for the theorems. As an approach to automated theorem proving over the real numbers, the scalability can outperform existing symbolic approaches. For instance, Tom Hales' Flyspeck project [6], [7] for the formalization of his proof the Kepler conjecture, requires proving hundreds of nonlinear real inequalities. We will demonstrate that we can automatically generate proofs for many of such formulas.

It is worth pointing out that after proof generation, proof checking is still a nontrivial problem because of the use of numerical procedures in the computation. Indeed, not all of the unsat answers that we have obtained can be proof checked. The challenge lies in validating basic axioms about nonlinear functions over the reals, which can be easily established by numerical algorithms (such as Newton iteration), but not symbolically. Ideally, we need to formalize most of the numerical algorithm in a δ-complete decision procedure to achieve full validation. We regard this as an interesting direction towards bridging the gap between numerical and symbolic methods in solving formulas over the real numbers.

We will describe our approach in the following steps:

1. We formalize the ICP algorithm in the framework of Abstract DPLL [13]. The similarity between ICP and SAT solving techniques has been explored in existing work [3]. With this formulation, the branch-and-prune framework is viewed as a transition system with a small set of transition rules. (Section II)

2. We use a simple first-order proof calculus $\mathbb{D}_A$, relativized to a set A of axioms over the reals, and show how to transform a run of the Abstract ICP to a proof in the system. (Section III)

3. We show how to validate the generated proofs using a stand-alone proof checker implementing simple rules and reliable interval arithmetic. The proof checker interacts with the solver in an abstraction refinement loop to obtain proof trees of sufficient detail (Section IV). In the end, we show experimental results towards the proving nonlinear lemmas in the Flyspeck project, in Section V.

Related Work.: Our work is closely related to several lines of research in the existing literature. For proving formulas with transcendental functions, MetiTarski [16], [1], [15] is the leading tool that reduces problems to polynomials and calls quantifier elimination procedures. Note that in MetiTarski, the polynomial problems are solved using external tools, without producing proofs. For problems with only polynomials, Bernstein polynomials are used in PVS for formal proofs [12]. Our approach aims to automatically produce complete formal proofs for formulas with transcendental functions. The iSAT solver [3] also contains strategies for certifying their answers in a different framework [10]. There are now several SMT solvers [8], [14] for formulas with nonlinear polynomials over the reals based on CAD with no proof-producing capacities, but a proof-producing algorithm is possible, as sketched in [11]. Proofs for correctness in general SMT solvers have been well studied, for instance in [18], which allows us to focus on the nonlinear theory solver in our framework.

II. A Formalization of Interval Constraint Propagation

Interval Constraint Propagation (ICP) [2] finds solutions of real constraints using the "branch-and-prune" method, com-

bining interval arithmetic and constraint propagation. The idea is to use interval extensions of functions to "prune" out sets of points that are not in the solution set and "branch" on intervals when such pruning can not be done, recursively until a small enough box that may contain a solution is found or inconsistency is observed. A high-level description of the decision version of ICP is given in Algorithm 1 [2], [4].

Algorithm 1 ICP($f_1, ..., f_m, B_0 = I_1^0 \times \cdots \times I_n^0, \delta$)

$S \leftarrow B_0$
while $S \neq \emptyset$ **do**
  $B \leftarrow S.\text{pop}()$
  while $\exists 1 \leq i \leq m, B \neq_\delta \text{Prune}(B, f_i)$ **do**
    $B \leftarrow \text{Prune}(B, f_i)$
  end while
  if $B \neq \emptyset$ **then**
    if $\exists 1 \leq i \leq m, |\sharp f_i(B)| \geq \delta$ **then**
      $\{B_1, B_2\} \leftarrow \text{Branch}(B, i)$
      $S.\text{push}(\{B_1, B_2\})$
    else
      return sat
    end if
  end if
end while
return unsat

Our task now is to formalize ICP algorithms so that we can extract symbolic proofs from its computation processes. Similar to Abstract DPLL, we represent ICP as a transition system, whose states consist of interval assignments and the real constraints to be solved. An interval I is any connected subset of $\mathbb{R}$ and we write $\mathbb{IR}$ to denote the set of all the intervals. We first formalize how ICP maintains interval assignments to a set of variables as follows:

Definition 1 (Interval Assignment Sequence)**.** *Let $x_1, ..., x_n$ be real variables. An* interval assignment sequence *over $\vec{x}$ is a sequence $(s_1, ..., s_m)$, where*

$$s_i \in \{(x_i \in I_j) : 1 \leq i \leq n, I_j \in \mathbb{IR}\} \cup \{(x_i \in I_j)^d : 1 \leq i \leq n, I_j \in \mathbb{IR}\}.$$

We write (S_1, S_2) to denote the concatenation of two sequences S_1 and S_2. The parentheses can be omitted when appropriate.

It will be clear later that when we write $(x \in I)^d$, it means an arbitrary choice on the value of x (called a d-assignment), which is consequently a backtrack point.

Remark 2. *ICP can maintain unions of intervals for variables. In principle this is not needed if we only consider the decision problem, which only searches for one solution and the components of a union can be tested sequentially. So we assume that only connected subsets of values are used here.*

Definition 3 (Box Domain)**.** *Let S be an interval assignment sequence over variables $x_1, ..., x_n$. The* box domain *associated with S is defined by*

$$\beta(S) = I_1 \times \cdots \times I_n,$$

where $I_i = \bigcap\{I : (x_i \in I)$ or $(x_i \in I)^d$ occurs in $S\}$. Also, we write $\beta(S)_i$ to denote I_i.

Definition 4 (ICP Transitions)**.** *Let $\vec{x} = (x_1, ..., x_n)$ be a vector of real variables. We write $c(\vec{x})$ to denote a constraint over $\mathbb{R}^n$, and S an interval assignment sequence over $\vec{x}$. Let $S \parallel c$ be the current state. We will always write $\beta(S_i) = I_i$ to denote the current interval assignment on variable x_i. We now define the following transition rules from $S \parallel c$ to another state.*

(Pruning): Let I_i^1 be a subset of I_i such that $\forall \vec{a} \in \beta(S, x_i \in I_i^1)$, $c(\vec{a})$ is false. Then, if we let I_i^2 be an interval satisfying $I_i \subseteq I_i^1 \cup I_i^2$, then

$$S \parallel c \quad \overset{p}{\Longrightarrow} \quad S, (x_i \in I_i^2) \parallel c$$

is called a pruning step.

(Branching): Let I_i^1 be a subset of I_i. Then

$$S \parallel c \quad \overset{br}{\Longrightarrow} \quad S, (x_i \in I_i^1)^d \parallel c,$$

is called a branching step.

(Backtracking): Let I_i^1 be a subset of I_i, such that $\forall \vec{a} \in \beta(S, x_i \in I_i^1, S')$, $c(\vec{a})$ is false. Let I_i^2 be an interval such that $I \subseteq I_i^1 \cup I_i^2$. If in addition, S' does not contain any d-assignment (of the form $(x \in I)^d$), then we can make a transition

$$S, (x_i \in I_i^1)^d, S' \parallel c \overset{bt}{\Longrightarrow} \quad S, (x_i \in I_i^2) \parallel c,$$

which is called a backtracking step.

(Failure): Suppose $\forall \vec{a} \in \beta(S)$, $c(\vec{a})$ is false, and there is no d-assignment in S. Then we can make the transition

$$S \parallel c \overset{f}{\Longrightarrow} \emptyset \parallel c$$

which is called a failure step.

Definition 5 (Abstract ICP)**.** *An n-dimensional ICP framework is a transition system*

$$\langle \mathbb{IR}^n, \mathcal{S}, \mathcal{C}, \Longrightarrow, \varepsilon \rangle$$

where $\mathcal{S}$ is the set of all interval assignment sequences over $\mathbb{IR}^n$, and $\mathcal{C}$ is any set of constraints over $\mathbb{R}^n$. A state is an element in $\mathcal{S} \parallel \mathcal{C}$. The transition rules $\Longrightarrow$: $\mathcal{S} \times \mathcal{C} \to \mathcal{S} \times \mathcal{C}$ are as specified in Definition 4. $\varepsilon \in \mathbb{Q}^+$ is an error bound. A run *of ICP is any sequence*

$$S_1 \parallel c, ..., S_k \parallel c,$$

where either S_k is $\emptyset$, or $S_k \neq \emptyset$ and $||\beta(S_k)|| < \varepsilon$.

Remark 6. *We have defined ICP in a general way, without enforcing conditions on the pruning operators, such as well-definedness. Thus, many invalid ICP runs can be generated. In this way, we treat ICP as a proof searching algorithm, and rely on the proof checkers to determine the correctness of an ICP run. In practice, of course, only "correct" ICP algorithms can provide proofs that can always be validated.*

Example 7. *Consider a constraint $c(x, y) : y = x \wedge y =$*

x^2, *and* $x \in [1.5, 2]$ *and* $y \in [1, 2]$ *are the initial interval assignment. A possible ICP run is:*

$$
\begin{aligned}
& x \in [1.5,2], y \in [1,2] \parallel c \\
\overset{br}{\Longrightarrow}\ & x \in [1.5,2], y \in [1,2], (x \in [1.7,2])^d \parallel c \\
\overset{bt}{\Longrightarrow}\ & x \in [1.5,2], y \in [1,2], x \in [1.5,1.7] \parallel c \\
& \textit{(backtracking, since } \forall \vec{a} \in [1.7,2] \times [1,2],\ c(\vec{a}) \textit{ is false,} \\
& \quad \textit{and } [1.5,2] \subseteq [1.5,1.7] \cup [1.5,2] \textit{ for } x) \\
\overset{p}{\Longrightarrow}\ & x \in [1.5,2], y \in [1,2], x \in [1.5,1.7], x \in [1.5,1.6] \parallel c \\
& \textit{(pruning, since } \forall \vec{a} \in [1.6,1.7] \times [1,2],\ c(\vec{a}) \textit{ is false)} \\
\overset{p}{\Longrightarrow}\ & x \in [1.5,2], y \in [1,2], x \in [1.5,1.7], \\
& \qquad\qquad x \in [1.5,1.6], x \in \emptyset \parallel c \\
& \textit{(pruning, since } \forall \vec{a} \in [1.5,1.6] \times [1,2],\ c(\vec{a}) \textit{ is false)} \\
\overset{f}{\Longrightarrow}\ & \emptyset || c \quad \textit{(since } \forall \vec{a} \in \emptyset \times [1,2], c(\vec{a}) \textit{ is false.)}
\end{aligned}
$$

III. Extracting Proofs from ICP Runs

A. First-Order Proofs of Unsatisfiability

We focus on the proof the unsatisfiability of conjunctions of theory atoms in the DPLL(T) framework, i.e., formulas of the form

$$\exists^{I_1} x_1 \cdots \exists^{I_n} x_n.\ \bigwedge_{i=1}^{m} f_i(x_1, ..., x_n) \sim 0$$

where $\sim \in \{=, \neq, >, \geq, <, \leq\}$. It is clear that once such proofs are obtained, the proof of unsatisfiability of Boolean combinations of the theory atoms can be obtained, by simply plugging them in the high level resolution proof. Also, it is important to note that the ICP algorithm solves *systems* of constraints, and it regards the conjunction $\bigwedge_{i=1}^{m} f_i(x_1, ..., x_n) \sim 0$ as one constraint $c(x_1, ..., x_n)$. Consequently, our task is now reduced to obtaining proofs for the validity of formulas of the form $\forall x_1 \cdots \forall x_n.(x_1 \in I_1 \wedge \cdots \wedge x_n \in I_n) \rightarrow \neg c(\vec{x})$, from the failure of ICP search for a solution to the original SMT formula $\exists \vec{x}.\vec{x} \in \vec{I} \wedge c(\vec{x})$.

We will construct a simple first-order proof calculus, and show how to transform ICP runs into proofs in the system.

Again, we consider formulas in a signature $\mathcal{L}_F = \langle <, \mathcal{F} \rangle$, where constants are considered as 0-ary functions in $\mathcal{F}$. When we write $x \in I$, where I denotes an interval, it is regarded as an abbreviation for their equivalent $\mathcal{L}_\mathcal{F}$-formula. Note that this means that I only uses $\mathcal{L}_\mathcal{F}$-terms as end-points. Also, as mentioned above, $c(\vec{x})$ abbreviates a conjunction of atomic formulas. We also allow the use of vectors in the formulas, writing $\vec{x} \in \vec{I}$ to denote $\bigwedge_i x_i \in I_i$.

Definition 8 (System $\mathbb{D}_A$)**.** *We define $\mathbb{D}_A$ to be the first-order proof system consisting of only the following two rules:*

$$\frac{\forall \vec{x}(\psi \rightarrow \varphi) \qquad \forall \vec{x}(\psi' \rightarrow \varphi)}{\forall \vec{x}(\psi \vee \psi' \rightarrow \varphi)}\ \vee I$$

$$\frac{\forall \vec{x}(\psi \rightarrow \varphi) \qquad \forall \vec{x}(\psi' \rightarrow \psi)}{\forall \vec{x}(\psi' \rightarrow \varphi)}\ \forall MP$$

and a set A of axioms of the following two types:

Interval Axioms:

$$\frac{}{\forall \vec{x}(\vec{x} \in \vec{I} \rightarrow \vec{x} \in \vec{I}_1 \vee \vec{x} \in \vec{I}_2)}\ IA$$

Constraint Axioms:

$$\frac{}{\forall \vec{x}(\vec{x} \in \vec{I} \rightarrow c(\vec{x}))}\ CA$$

Derivations in $\mathbb{D}_A$ are as standardly defined, as natural deductions following these rules. Clearly, the two first-order rules are valid. Thus, if all the axioms in A are valid, then the system only produces valid formulas over $\mathbb{R}$.

Proposition 9 (Soundness)**.** *If $\mathbb{D}_A \vdash \varphi$ and $\mathbb{R} \models \bigwedge A$, then $\mathbb{R} \models \varphi$.*

Remark 10. *Clearly, the constraint axioms are the most nontrivial part. They are the basic facts of real functions that a numerical procedure relies on, usually concerning the range of functions within a small interval. The interval axioms are sometimes not trivial either (for instance, compare intervals ending with e^π and π^e respectively). Proof-checking involves validation of these axioms, which we discuss in Section IV.*

We now describe the construction of proof trees from ICP runs, which will be represented as labeled binary trees. A labeled binary tree is defined as a tuple $T = \langle V, V_L, \Sigma, \delta, \sigma \rangle$. Here, $V = \{v_0, ..., v_k\}$, is a finite set of nodes, where $v_0 \in V$ always denotes the root node. V_L is the set of leaf nodes in V. Σ is a set of labels, which in our case is the set of $\mathcal{L}_\mathcal{F}$-formulas. $\delta :\subseteq V \times \{l, r\} \rightarrow V$ is a partial mapping from a node to its descendant nodes, where $\delta(v, l)$ and $\delta(v, r)$ denote the left and right descendant nodes, respectively. $\sigma :\subseteq V \rightarrow \Sigma$ is a labeling function that maps each node $v \in V$ to a formula $\sigma(v) \in \Sigma$. In addition, the edges in the tree can be labeled as well, through a function $\tau : V \times V \rightarrow \Omega$ where Ω is a set of edge-labels.

1) Tree Generation: Let an ICP run be

$$S_0 \parallel c \overset{t_1}{\Longrightarrow} \cdots \overset{t_m}{\Longrightarrow} S_m \parallel c,$$

such that the ending transition t_m is a failure step, i.e., $S_m = \emptyset$. We now define the procedure by defining the functions δ and V_L through induction on s_i. The edges can be labeled naturally with Ω = {$\vee$I, $\forall$M, IA, CA}.

Case $i = 0$.: We label the root node v_0 by

$$\sigma(v_0) := \forall \vec{x}(\vec{x} \in \beta(S_0) \rightarrow \neg c).$$

Let $V_L^0 = \{v_0\}$ denote the current collection of leaf nodes. Note that this formula is the negation of the input SMT formula.

Case $i = k + 1$ ($1 < k \leq m$). : Suppose V_L^k and σ have been defined for $s_1, ..., s_k$. Write $s_k = S_k \parallel c$ and $s_{k+1} = S_{k+1} \parallel c$. Now we split the cases on the type of the step t from s_k to s_{k+1} as follows. Again, we use the convention that $\beta(S)_i = I_i$ denotes the current interval assignment on a variable x_i.

(Pruning Case): Suppose $s_k \Longrightarrow s_{k+1}$ is a pruning step. That is,

$$S_k \parallel c \overset{p}{\Longrightarrow} S_k, (x_i \in I_i^2) \parallel c,$$

where $I_i \subseteq I_i^1 \cup I_i^2$ and $\forall \vec{a} \in \beta(S_k, x_i \in I_i^1)$, $c(\vec{a})$ is false. If we write

$$\vec{I}_1 = \beta(S_k, (x_i \in I_i^1)), \vec{I}_2 = \beta(S_k, (x_i \in I_i^2)), \text{ and } \vec{I} = \beta(S_k),$$

then this step corresponds to the sub-tree as shown in Fig. 1, Case A.

Formally, the sub-tree is added as follows. Let $v \in V_L^k$ be an existing leaf node that is labeled by the formula corresponding to $S_k \parallel c$; namely,

$$\sigma(v) = \forall \vec{x}(\vec{x} \in \vec{I} \rightarrow \neg c).$$

(We will inductively prove that such a node exists.) We then define

$$\begin{aligned}
\delta(v, l) &= v_{k+1}^1, \sigma(v_{k+1}^1) = \forall \vec{x}((\vec{x} \in \vec{I}_1 \vee \vec{x} \in \vec{I}_2) \rightarrow \neg c); \\
\delta(v, r) &= v_{k+1}^2, \sigma(v_{k+1}^2) = \forall \vec{x}(\vec{x} \in \vec{I} \rightarrow (\vec{x} \in \vec{I}_1 \vee \vec{x} \in \vec{I}_2)); \\
\delta(v_{k+1}^1, l) &= v_{k+1}^3, \sigma(v_{k+1}^3) = \forall \vec{x}(\vec{x} \in \vec{I}_2 \rightarrow \neg c) \\
\delta(v_{k+1}^1, r) &= v_{k+1}^4, \sigma(v_{k+1}^4) = \forall \vec{x}(\vec{x} \in \vec{I}_1 \rightarrow \neg c)
\end{aligned}$$

and set $V_L^{k+1} = (V_L^k \setminus \{v\}) \cup \{v_{k+1}^3\}$.

(Branching Case): Suppose $s_k \Longrightarrow s_{k+1}$ is a branching step. That is,

$$S_k \parallel c \quad \stackrel{br}{\Longrightarrow} \quad S_k, (x_i \in I_i^1)^d \parallel c,$$

under the condition that $I_i^1 \subseteq I_i$. If we write

$$\vec{I}_1 = \beta(S, (x_i \in I_i^1)), \vec{I}_2 = \beta(S, (x_i \in I_i^2)), \text{ and } \vec{I} = \beta(S),$$

where $I \subseteq I_i^1 \cup I_2$, then this step corresponds to the sub-tree as shown in Fig. 1, Case B. Formally it is defined as follows. Again, let $v \in V_L^k$ be a leaf node such that $\sigma(v) = \forall \vec{x}(\vec{x} \in \vec{I} \rightarrow \neg c)$. We then define

$$\begin{aligned}
\delta(v, l) &= v_{k+1}^1, \sigma(v_{k+1}^1) = \forall \vec{x}(\vec{x} \in \vec{I}_1 \vee \vec{x} \in \vec{I}_2 \rightarrow \neg c); \\
\delta(v, r) &= v_{k+1}^2, \sigma(v_{k+1}^2) = \forall \vec{x}(\vec{x} \in \vec{I} \rightarrow (\vec{x} \in \vec{I}_1 \vee \vec{x} \in \vec{I}_2)); \\
\delta(v_{k+1}^1, l) &= v_{k+1}^3, \sigma(v_{k+1}^3) = \forall \vec{x}(\vec{x} \in \vec{I}_1 \rightarrow \neg c) \\
\delta(v_{k+1}^1, r) &= v_{k+1}^4, \sigma(v_{k+1}^4) = \forall \vec{x}(\vec{x} \in \vec{I}_2 \rightarrow \neg c)
\end{aligned}$$

and set $V_L^{k+1} = (V_L^k \setminus \{v\}) \cup \{v_{k+1}^3, v_{k+1}^4\}$.

(Backtracking Case): Suppose $s_k \Longrightarrow s_{k+1}$ is a branching step. That is,

$$S_{k'}, (x_i \in I_i^1)^d, S' \parallel c \quad \stackrel{bt}{\Longrightarrow} \quad S_{k'}, (x_i \in I_i^2) \parallel c,$$

when $\forall a \in \beta(S, (x_i \in I_i^1)^d, S')$, $c(\vec{a})$ is false, and $I_i \subseteq I_i^2 \cup I_i^1$, where $I_i = \beta(S_{k'})_i$. $S_{k'}$ is a previous interval assignment sequence, with $k' < k$. If we write

$$\vec{I}_1 = \beta(S, (x_i \in I_i^1)), \vec{I}_2 = \beta(S, (x_i \in I_i^2), \text{ and } \vec{I} = \beta(S_{k'}),$$

then this step corresponds to the sub-tree as shown in Fig. 1, Case C. Formally, we simply set $V_L^{k+1} = V_L^k$, and do not update the labeling.

(Fail Case): Suppose it is a failure step. That is,

$$S \parallel c \quad \stackrel{f}{\Longrightarrow} \quad \emptyset \parallel c$$

when $\forall \vec{a} \in \beta(S)$, $c(\vec{a})$ is false and S has no d-assignments. Let $\vec{I} = \beta(S)$. This step corresponds to

$$\frac{}{\forall \vec{x}(\vec{x} \in \vec{I}) \rightarrow \neg c}\text{FA}$$

We set $V_L^{k+1} = V_L^k \setminus \{v\}$ and do not update σ.

Complete tree.: In all, after all the steps in the ICP run are followed, the tree that we construct is $T = \langle V, V_L^m, \Sigma, \delta, \sigma \rangle$. The axiom set is given by

$$A = \{\sigma(v) : v \in V_L^m\}.$$

It is easy to see that T is a valid proof tree in $\mathbb{D}_A$:

Proposition 11. *For every ICP run ending with $\emptyset \parallel c$, the tree construction procedure above produces a valid natural deduction tree in $\mathbb{D}_A$. The size of the proofs is linear in the computation steps.*

Proof: It is clear that each proof step, as represented by the subtree created in each case, is a valid natural deduction step in $\mathbb{D}_A$. We only need to show that the tree can be constructed. For this, we need to show that for each step $S_k \parallel c \stackrel{t}{\Longrightarrow} S_{k+1} \parallel c$, where S_{k+1} is not $\emptyset$, it is always the case that $S_k \parallel c$ labels a leaf node in the tree constructed so far. When $k = 0$, this is the case since $V_L^0 = \{v_0\}$. Now suppose $S_k \parallel c$ labels a leaf node. If t is a pruning step, then $\forall \vec{x}(\vec{x} \in \vec{I}_2 \rightarrow \neg c)$ labels v_{k+1}^3, which is added in V_L^{k+1}. The same applies to the other branching and backtracking. Finally, the step $S_{m-1}||c \Longrightarrow \emptyset||c$ corresponds to closing the last leaf labelled by $\forall \vec{x}(\vec{x} \in \vec{I} \rightarrow \neg c)$. ■

Again, once the proof tree is constructed, the details of the ICP algorithm no longer matters. The only rules involved are the two first-order rules in $\mathbb{D}_A$. Following relative soundness of the system, to establish validity of the formula, now we only need to validate the axiom set A.

IV. Validating the Proofs

A. Validating the Axioms

There are two types of axioms that we allow in the proofs constructed from ICP runs: interval axioms and constraint axioms. To validate such axioms, we still need numerical computations. The difference is that the proof checker only needs to implement simple interval computation that can be validated through stand-alone arbitrary-precision or symbolic computation. Note that the validation of the axioms can fail when the solver correctly returns unsat, if the solver uses complex numerical heuristics that can not be verified by reliable numerical computation. In practice, we ensure the correctness of the proof checker first, and use an abstraction refinement loop that allows the proof checker to ask for more detailed proofs from the solver.

The interval axioms do not contain any real functions, and are of the form $\forall \vec{x}(x \in I_1 \vee x \in I_2 \rightarrow x \in I)$. We only need to check that I is a subset of $I_1 \cup I_2$ by comparing the end points of the intervals, which is an easy numerical task.

The constraint axioms are of the form $\forall x(\vec{x} \in \vec{I} \rightarrow c(\vec{x}))$, and can only be verified by considering the functions that occur in c. Although they are of the same form as the formulas

A. Pruning Case:

$$\dfrac{\dfrac{\dfrac{\vdots}{\forall\vec{x}(\vec{x}\in\vec{I}_2\to\neg c)}\quad\dfrac{}{\forall\vec{x}(\vec{x}\in\vec{I}_1\to\neg c)}\text{CA}}{\forall x((\vec{x}\in\vec{I}_1\vee\vec{x}\in\vec{I}_2)\to\neg c)}\vee\text{I}\qquad\dfrac{}{\forall x(x\in I_i\to(x\in I_i^1\vee x\in I_i^2))}\text{IA}}{\forall\vec{x}(\vec{x}\in\vec{I}\to\neg c)}\forall\text{MP}$$

B. Branching Case:

$$\dfrac{\dfrac{\dfrac{\vdots}{\forall\vec{x}(\vec{x}\in\vec{I}_1\to\neg c)}\quad\dfrac{\vdots}{\forall\vec{x}(\vec{x}\in\vec{I}_2\to\neg c)}}{\forall x(\vec{x}\in\vec{I}_1\vee\vec{x}\in\vec{I}_2\to\neg c)}\vee\text{I}\qquad\dfrac{}{\forall x(x\in I_i\to(x\in I_i^1\vee x\in I_i^2))}\text{IA}}{\forall\vec{x}(\vec{x}\in\vec{I}\to\neg c)}\forall\text{MP}$$

C. Backtracking Case:

$$\dfrac{\dfrac{\dfrac{\dfrac{}{\forall\vec{x}(\vec{x}\in\beta(S_{k'},(x\in I_i^1)^d,S')\to\neg c)}\text{CA}}{\forall\vec{x}(\vec{x}\in\vec{I}_1\to\neg c)}\quad\dfrac{\vdots}{\forall\vec{x}(\vec{x}\in\vec{I}_2\to\neg c)}}{\forall x(\vec{x}\in\vec{I}_1\vee\vec{x}\in\vec{I}_2\to\neg c)}\vee\text{I}\qquad\vdots}{\forall\vec{x}(\vec{x}\in\vec{I}\to\neg c)}\forall\text{MP}$$

Fig. 1: Proof Trees

we solve, these axioms should contain evident properties of the functions involved, usually on small intervals. Such facts can be verified using reliable interval computations, for instance as follows.

Definition 12 (Interval Extensions)**.** *Let $f:\mathbb{R}^n\to\mathbb{R}$ be a real function. An interval function $F:\mathbb{IR}^n\to\mathbb{IR}$ is a function that satisfies:*

$$\forall I\in dom(F),\{f(x):x\in I\}\subseteq F(I).$$

A simple example of interval extensions is the *natural interval extension* for arithmetic operations, based on computations of functions on the end points of intervals. It is obvious that:

Proposition 13. *Let F be an interval extension of f, and $I\subseteq dom(f)$. If $F(I)\subseteq A$, then $\forall x(x\in I\to f(x)\in A)$.*

Thus, the axioms are validated if we can verify that they are consistent with all the interval extensions.

Example 14. *The second pruning step in Example 7 generates an axiom*

$$\forall x\forall y(x\in[1.7,2]\wedge y\in[1,2]\to\neg(y=x^2)\vee\neg(y=x))$$

This can be easily validated through the natural interval extension of $(y-x^2)$, which is $[1,2]-[1.7,2]^2=[-3,-0.89]$ and does not contain 0.

B. Taylor Proofs

Suppose we want to verify the inequality $f(x_1,...,x_n)>0$ on a domain $\vec{x}\in D=I_1\times\cdots\times I_n$. Using the multivariate mean value theorem, we have that for any $\vec{a},\vec{b}\in D$

$$f(\vec{b})-f(\vec{a})=\nabla f(\xi)\cdot(\vec{b}-\vec{a})=\sum_i\frac{\partial f}{\partial x_i}(\xi)\cdot(b_i-a_i)$$

for some $\xi\in D$. Thus, we can bound $f(x)$ on D by computing the interval bound on the function

$$f(\vec{a})+\sum_i\left(\sharp\left(\frac{\partial f}{\partial x_i}\right)(D)\right)\cdot D|_{x_i}$$

where $\sharp(\cdot)$ denotes interval extension, and $f(\vec{a})$ is on the boundary of D.

Example 15. *$f(x_1,x_2)=x_1^2+x_2^2$ on domain $(x_1,x_2)\in[0,1]\times[0,1]$. We have $\partial f/\partial x_1=2x_1\in[0,2]$ and $\partial f/\partial x_2\in[0,2]$. Thus*

$$f(\vec{x})\in\sum_{i=1,2}[0,2]\cdot(1-0)+0=[0,4].$$

C. The Branch and Prove Loop

In practice, ICP usually implements complicated heuristics that are more powerful than what can be verified through direct interval arithmetic. A practical approach first is to use an abstraction refinement loop that allows the proof checker to ask the solver for proof traces of the right amount of details. We sketch the procedures in Algorithm 2 and Algorithm 3.

When we fail to prove an axiom through simple interval arithmetic, the proof checker generates new subproblems that are returned to the solver. At this stage, the axioms become the new theorems to be proved. This is an abstraction refinement procedure. Algorithm 3 illustrates the procedure. By executing the loop, we may obtain proof trees that contain more and

Algorithm 2 ProofCheck

procedure PROOFCHECK(p, δ)
  if MATCH(p, Axiom($\forall x(\vec{x} \in \vec{I} \to c(\vec{x}))$)) **then**
    if $\sharp c(\vec{I})$ **then** ▷ use interval arithmetic, taylor extension..
      return $\emptyset$
    else
      $(\vec{I_1}, \vec{I_2}) \leftarrow \text{Split}(\vec{I})$
      $(\delta_1, \delta_2) \leftarrow (\min(\delta, \frac{1}{4}||I_1||), \min(\delta, \frac{1}{4}||I_2||))$
      return $\{(\forall x(\vec{x} \in \vec{I_1} \to c(\vec{x})), \delta_1), (\forall x(\vec{x} \in \vec{I_2} \to c(\vec{x})), \delta_2)\}$
    end if
  else if MATCH(p, Branch($p_1, p_2, \vec{I}$)) **then**
    $U_1 \leftarrow$ PROOFCHECK(p_1, δ)
    $U_2 \leftarrow$ PROOFCHECK(p_2, δ)
    if $\vec{I} \not\subseteq (\text{dom}(p_1) \cup \text{dom}(p_2))$ **then**
      return Error
    else
      return $U_1 \cup U_2$
    end if
  end if
end procedure

Algorithm 3 Branch-and-Prove

procedure BRANCH-AND-PROVE(p, δ)
  $U \leftarrow$ PROOFCHECK(p, δ)
  if $U \neq \emptyset$ **then**
    for all $(a, \delta') \in U$ **do**
      $p' \leftarrow$ SOLVE(a, δ')
      BRANCH-AND-PROVE(p', δ')
    end for
  end if
end procedure

more detailed steps. There are two ways that the prover can generate the subproblems, branching on a variable in the formula or using a smaller δ. Note that under the condition that the pruning operators in the solver is well-defined, both procedures never change the unsat result. The branching may give exponentially many new problems; while the δ-change does not give new problems, but may exponentially slow down the solver in each round. In practice we observe that such a refinement loop is very useful, as we will show in the experiments.

V. EXPERIMENTS

We implemented the proof generation capacity into our open-source solver dReal[1]. All the experiments below are performed on a machine of with a 32-core 2.0GHz Intel Xeon E5-2600 Processor and 64GB of RAM. The benchmarks and full tables of experiment statistics are also available on the tool page.

A main set of benchmarks that we studied is from the Flyspeck project [6], [17], which aims at a fully formalized proof of the Kepler conjecture. As lemmas for the proof, hundreds of nonlinear real inequalities need to be verified. Although the formulas usually contain only around ten variables, they contain a huge number of nonlinear arithmetic operations and trigonometric functions, and are mathematically challenging. In the original proof, Hales implemented procedures that combine linear programming and interval arithmetic to establish all these formulas, but the algorithms are formally verify. In fact, the formal verification of these nonlinear inequalities is the last main piece of work needed to complete the full project. Without any particular optimization on ICP, we have observed promising results. Out of 916 nonlinear formulas in the Flyspeck project repository, the solver returns unsat for 107 of them with a timeout of 5 minute each, and a precision $\delta = 10^{-3}$. Out of these formulas, we automatically generated and validated the proofs for 72 instances. The proof traces of these formulas can be very large; for instance, we proved one with more than 2M lines in the proof (54MB file). In Table II, we list some of the representative benchmarks to show scalability. Many of these formulas are highly nonlinear, for instance the formula encoded in 760.smt2 is following one

$$\forall \vec{x} \in [4.0, 6.3504]^5 \left(2\arctan\left(\frac{\Delta_2(\vec{x})}{\sqrt{\Delta_1(\vec{x}) + \Delta_2(\vec{x})^2} + \sqrt{\Delta_1(\vec{x})}}\right) - 0.458(\sqrt{x_2}+\sqrt{x_3}+\sqrt{x_4}+\sqrt{x_5}) + 0.342\sqrt{x_1} + 3.319204\right) < 0.0$$

where

$$\begin{aligned}
\Delta_1(\vec{x}) &= 4x_1(8x_1(-x_1 + x_2 + x_3 + x_4 + x_5 - 8) \\
&\quad + x_2x_5(x_1 - x_2 + x_3 + x_4 - x_5 + 8 \\
&\quad + x_3x_4(x_1 + x_2 - x_3 - x_4 + x_5 + 8) + 8x_2x_3 \\
&\quad - x_1x_3x_5 - x_1x_2x_4 - 8x_4x_5)) \\
\Delta_2(\vec{x}) &= x_2x_5 - x_2x_3 + x_3x_4 - x_4x_5 + x_1^2 - x_1x_2 \\
&\quad - x_1x_3 - x_1x_4 - x_1x_5
\end{aligned}$$

On the other hand, as mentioned above, we fail to establish about the proofs of unsatisfiability of about 30 instances. Table 2 shows some of these instances. They typically generate proofs that are large in size, or that the branch-and-prove loop has to generate too many sub-instances such that the proof checking can not terminate.

VI. CONCLUSION

We presented our approach for extracting formal proofs from a numerically-driven decision procedure in the DPLL$\langle$ICP$\rangle$ framework. We formalized the ICP algorithm, and showed how to validate proof trees from the unsat answers. A main focus for our tool is to prove nonlinear lemmas in the Flyspeck project, and we have observed promising experimental results. We believe the approach can be combined with existing symbolic methods, and is a first step towards a framework that bridges the gap between symbolic and numerical approaches. Further work would involve formalization of numerical algorithms, proof abstractions, local heuristics, and an implementation of our proof checker in standard proof assistants.

[1] http://dreal.cs.cmu.edu

ID	#Var	#Arith	Nonlinear	Time_S	Proof Size	#Sub	#Axiom	Time_PC
461	6	36	poly	1.740	2145155	2	17442	203.886
789	6	86	atan2,sqrt	1.640	350329	2	2464	128.077
792	6	828	atan2,sqrt	0.400	19837	2	118	113.004
745	6	36	poly	0.750	677580	2	5222	59.865
785	6	80	atan2,sqrt	0.470	63388	2	526	26.450
760	6	2767	atan2,sqrt	0.140	711	2	5	21.089
820	6	95	atan2,sqrt	0.080	9134	2	54	14.703
815	6	95	atan2,sqrt	0.330	41954	2	279	14.703
814	6	95	atan2,sqrt	0.350	42102	2	278	14.703
816	6	96	atan2,sqrt	0.110	12195	2	92	4.994
817	6	96	atan2,sqrt	0.090	11792	2	93	4.993
784	6	80	atan2,sqrt	0.060	7203	2	56	3.595
781	6	86	atan2,sqrt	0.060	7481	2	45	2.657
793	6	834	atan2,sqrt	0.020	18	1	1	1.855
796	6	834	atan2,sqrt	0.010	18	1	1	1.710
752	6	17	poly	0.080	46360	2	277	1.709
783	6	825	atan2,sqrt	0.020	93	1	1	1.549
779	6	201	atan2,sqrt	0.010	10	1	1	0.705
867	6	17	poly	0.040	25820	2	147	0.683
742	6	55	acos,atan2,sqrt	0.001	7	1	1	0.299
508	6	53	acos,sqrt	0.001	8	1	1	0.286
507	6	29	acos,sqrt	0.001	8	1	1	0.278
744	6	24	asin,cos,sin	0.001	8	1	1	0.275

TABLE I: Experimental results (Proved instances): ID = Problem ID, #Var = Number of variables, #Arith = Number of arithmetic operators, Nonlinear = Nonlinear operators occurred in problem, Proof Size = Number of lines of the proof, TIME_S = Solving time in seconds, #Sub = Number of subproblems generated by proof checking, #Axiom = Number of proved axioms, TIME_PC = Proof-checking time in seconds.

ID	#Var	#Arith	Nonlinear	Time_S	Proof Size	#Sub
260.smt2	6	90	poly	5.030	6281203	1
866.smt2	6	38	sqrt	0.390	543061	21476
775.smt2	6	2765	atan2,sqrt	4.040	130253	2
764.smt2	6	2767	atan2,sqrt	1.700	49657	2
762.smt2	6	2767	atan2,sqrt	2.040	42238	2
484.smt2	6	1835	acos,atan2,sqrt	0.060	16	1
485.smt2	6	1961	acos,atan2,sqrt	0.070	16	1
498.smt2	6	573	acos,matan,sqrt	0.010	11	8191

TABLE II: Experimental results (Unproved instances, Timeout = 300 sec): ID = Problem ID, #Var = Number of variables, #Arith = Number of arithmetic operators, Nonlinear = Nonlinear operators occurred in problem, Proof Size = Number of lines of the proof, TIME_S = Solving time in seconds, #Sub = Number of subproblems generated by proof checking,

REFERENCES

[1] B. Akbarpour and L. C. Paulson. MetiTarski: An automatic theorem prover for real-valued special functions. *J. Autom. Reasoning*, 44(3):175–205, 2010.

[2] F. Benhamou and L. Granvilliers. Continuous and interval constraints. In F. Rossi, P. van Beek, and T. Walsh, editors, *Handbook of Constraint Programming*, chapter 16. Elsevier, 2006.

[3] M. Fränzle, C. Herde, T. Teige, S. Ratschan, and T. Schubert. Efficient solving of large non-linear arithmetic constraint systems with complex boolean structure. *JSAT*, 1(3-4):209–236, 2007.

[4] S. Gao, J. Avigad, and E. M. Clarke. Delta-complete decision procedures for satisfiability over the reals. In Gramlich et al. [5], pages 286–300.

[5] B. Gramlich, D. Miller, and U. Sattler, editors. *Automated Reasoning - 6th International Joint Conference, IJCAR 2012, Manchester, UK, June 26-29, 2012. Proceedings*, volume 7364 of *Lecture Notes in Computer Science*. Springer, 2012.

[6] T. C. Hales. Introduction to the flyspeck project. In T. Coquand, H. Lombardi, and M.-F. Roy, editors, *Mathematics, Algorithms, Proofs*, volume 05021 of *Dagstuhl Seminar Proceedings*. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany, 2005.

[7] T. C. Hales, J. Harrison, S. McLaughlin, T. Nipkow, S. Obua, and R. Zumkeller. A revision of the proof of the kepler conjecture. *Discrete & Computational Geometry*, 44(1):1–34, 2010.

[8] D. Jovanovic and L. M. de Moura. Solving non-linear arithmetic. In Gramlich et al. [5], pages 339–354.

[9] S. Kong, S. Gao, and E. Clarke. Floating-point bugs in the embedded gnu c library, 2013. CMU SCS Technical Report CMU-CS-13-130.

[10] S. Kupferschmid, B. Becker, T. Teige, and M. Fränzle. Proof certificates

and non-linear arithmetic constraints. In R. Kraemer, A. Pawlak, A. Steininger, M. Schölzel, J. Raik, and H. T. Vierhaus, editors, *DDECS*, pages 429–434. IEEE, 2011.

[11] S. McLaughlin and J. Harrison. A proof-producing decision procedure for real arithmetic. In *CADE*, pages 295–314, 2005.

[12] C. Muñoz and A. Narkawicz. Formalization of a representation of Bernstein polynomials and applications to global optimization. *Journal of Automated Reasoning*, 2012. Accepted for publication.

[13] R. Nieuwenhuis, A. Oliveras, and C. Tinelli. Solving sat and sat modulo theories: From an abstract davis–putnam–logemann–loveland procedure to dpll(t). *J. ACM*, 53(6):937–977, 2006.

[14] G. O. Passmore and P. B. Jackson. Combined decision techniques for the existential theory of the reals. In J. Carette, L. Dixon, C. S. Coen, and S. M. Watt, editors, *Calculemus/MKM*, volume 5625 of *Lecture Notes in Computer Science*, pages 122–137. Springer, 2009.

[15] G. O. Passmore, L. C. Paulson, and L. M. de Moura. Real algebraic strategies for metiTarski proofs. In J. Jeuring, J. A. Campbell, J. Carette, G. D. Reis, P. Sojka, M. Wenzel, and V. Sorge, editors, *AISC/MKM/-Calculemus*, volume 7362 of *Lecture Notes in Computer Science*, pages 358–370. Springer, 2012.

[16] L. C. Paulson. MetiTarski: Past and future. In L. Beringer and A. P. Felty, editors, *ITP*, volume 7406 of *Lecture Notes in Computer Science*, pages 1–10. Springer, 2012.

[17] A. Solovyev and T. C. Hales. Formal verification of nonlinear inequalities with taylor interval approximations. In G. Brat, N. Rungta, and A. Venet, editors, *NASA Formal Methods*, volume 7871 of *Lecture Notes in Computer Science*, pages 383–397. Springer, 2013.

[18] A. Stump, D. Oe, A. Reynolds, L. Hadarean, and C. Tinelli. Smt proof checking using a logical framework. *Formal Methods in System Design*, 42(1):91–118, 2013.